

*STOCK OPTIMIZING: MAXIMIZING REINFORCERS
PER SESSION ON A VARIABLE-INTERVAL
SCHEDULE*

ALAN SILBERBERG, RICHARD BAUMAN, AND STEVEN HURSH

THE AMERICAN UNIVERSITY AND WALTER REED ARMY INSTITUTE OF RESEARCH

In Experiment 1, 2 monkeys earned their daily food ration by pressing a key that delivered food according to a variable-interval 3-min schedule. In Phases 1 and 4, sessions ended after 3 hr. In Phases 2 and 3, sessions ended after a fixed number of responses that reduced food intake and body weights from levels during Phases 1 and 4. Monkeys responded at higher rates and emitted more responses per food delivery when the food earned in a session was reduced. In Experiment 2, monkeys earned their daily food ration by depositing tokens into the response panel. Deposits delivered food according to a variable-interval 3-min schedule. When the token supply was unlimited (Phases 1, 3, and 5), sessions ended after 3 hr. In Phases 2 and 4, sessions ended after 150 tokens were deposited, resulting in a decrease in food intake and body weight. Both monkeys responded at lower rates and emitted fewer responses per food delivery when the food earned in a session was reduced. Experiment 1's results are consistent with a strength account, according to which the phases that reduced body weights increased food's value and therefore increased subjects' response rates. The results of Experiment 2 are consistent with an optimizing strategy, because lowering response rates when food is restricted defends body weight on variable-interval schedules. These contrasting results may be attributed to the discriminability of the contingency between response number and the end of a session being greater in Experiment 2 than in Experiment 1. In consequence, subjects lowered their response rates in order to increase the number of reinforcers per session (stock optimizing).

Key words: economics, stock optimizing, income maximizing, variable-interval schedule, tokens, key press, token deposit, monkeys

Some behavior analysts have tried to determine whether behavior in nonhumans can be described in terms of optimality-based decision rules similar to those used in economics (e.g., Rachlin, Battalio, Kagel, & Green, 1981). Toward this end, researchers have created analogies in the operant chamber to human work, pay, and choice in a marketplace. Then they have tested whether nonhuman behavior can be described by decision rules similar to those used to account for the consumer behavior of humans. To date, several correspondences have been noted between decisions in the human economy and schedule-maintained behavior in nonhumans (for illustrations, see Lea, Tarpy,

& Webley, 1987). This suggests that common processes may underlie economic and operant behavior.

If common decision rules account for behavior in the marketplace and in the operant chamber, whose rules—those of the behavior analyst or those of the microeconomist—should we extend into the other's domain? This question has various answers. Some have argued that the optimality assumption of microeconomics can be successfully extended to explain reinforcement-schedule-maintained behavior in nonhumans (e.g., Rachlin et al., 1981), whereas others have claimed that behavior is often suboptimal in both the laboratory and the human economy, and can be better explained by behavior-analytic notions such as response strength (Vaughan & Miller, 1984) or melioration (Herrnstein, 1990).

Vaughan and Miller's (1984) study is illustrative of those claiming to demonstrate that an adequate model of behavior cannot be based on an optimality assumption. In Experiment 1 of their study, pigeons maintained at 80% of free-feeding weights responded on a variable-interval (VI) schedule specially constructed so that the reinforcement rate it provided was inversely related to response rate. Pigeons re-

This research was supported by a National Science Foundation grant to The American University. In conducting the research described in this report, the investigators adhered to the *Guide for the Care and Use of Laboratory Animals*, DHEW publication (NIH) 73-23, as prepared by the Institute of Laboratory Animal Resources, National Research Council. The views of the authors do not purport to reflect the position of the Department of the Army or the Department of Defense (para 4-3, AR 360-5). The authors acknowledge the help of N. Berendzen in training monkeys to deposit tokens. Address correspondence to A. Silberberg, Department of Psychology, The American University, Washington, DC 20016 (e-mail network address: asilber@auvm.american.edu).

sponded at high rates despite the fact that these rates reduced reinforcement rates to low levels. The birds' failure to maximize reinforcement rate might indicate that nonhumans maximize dimensions of reinforcement other than rate. Two likely candidates are (a) maximizing reinforcer immediacy (what we call *present-value maximizing*; see Silberberg, Warren-Boulton, & Asano, 1988) and (b) maximizing the amount of a good earned per day (*income maximizing*; see Sakagami, Hursh, Christensen, & Silberberg, 1989; Shurtleff & Silberberg, 1990).

The proposal that subjects maximize present value is based on the premise that nonhumans have a strong preference for short delays between responses and reinforcement. The tendency to maximize reinforcer immediacy can account not only for the "irrationally" high response rates obtained by Vaughan and Miller (1984) but also for other schedule-based behavioral inefficiencies, such as the occurrence of many unreinforced responses on VI and differential-reinforcement-of-low-rate schedules and the prematurity of most responding on fixed-interval schedules (see Silberberg *et al.*, 1988).

The proposal that subjects maximize income suggests that nonhumans respond so as to maximize the amount of a good earned in a day. Expressed in terms of food, this translates into ensuring at a minimum that daily caloric consumption equals daily energy expenditure (e.g., see Collier, Hirsch, & Kanarek, 1977; Shurtleff & Silberberg, 1990). In Vaughan and Miller's (1984) study, income maximizing (defense of food intake) was not possible because daily food consumption was independent of behavior: The birds' weights were maintained at a fixed level by supplemental feedings, if needed, regardless of performance during sessions. In consequence, we believe behavior was characterized solely by the alternate, present-value maximizing rule.

Given the opportunity, nonhumans may maximize income. The plausibility of this claim might be questioned in light of studies that seem to show that nonhumans often fail to defend food intake when such a strategy reduces immediacy, rate, or likelihood of a reinforcer between alternatives or conditions (e.g., Caraco, Martindale, & Whittam, 1980; Hamm & Shettleworth, 1987; Herrnstein & Heyman, 1979; Logue, Chavarro, Rachlin, & Reeder,

1988; Mazur, 1988; Timberlake, 1984; Timberlake, Gawley, & Lucas, 1987). Appearances notwithstanding, these studies do not contradict income-maximizing predictions because in none was daily food consumption allowed to vary as a function of behavior.

The distinction we draw between maximizing reinforcement rate and maximizing income mirrors the difference economists draw between income and wealth. They define the former in terms of money *flows* and the latter in terms of money *stocks* (Samuelson & Nordhaus, 1985, p. 564). For the economist, the flow-stock difference translates into different decision rules, depending on whether the phenomenon being modeled is attributed to an income or a wealth effect. In terms of this distinction, our use of the term *income maximizing* in the study of Shurtleff and Silberberg (1990), to refer to the goal of optimizing the total amount of reinforcement obtained in a session, seems ill-advised. This report rectifies this problem by using hereafter the term *stock optimizing* instead of *income maximizing*.

EXPERIMENT 1

In Vaughan and Miller's (1984) procedure, pigeons did not maximize the flow of reinforcement, and whether they would have maximized reinforcement stocks was not evaluated because daily food intake was kept constant and independent of behavior. Our first experiment sought evidence of stock optimizing in responding of monkeys on a VI schedule when all food was earned within the session. The "money" these monkeys spent was a depression of a response disk, and the commodity they "bought" was food pellets. Their wealth was manipulated across conditions by ending sessions after a specified number of responses. When the session terminated after many responses (in this circumstance, they were rich), a monkey could ignore the high price (responses per reinforcer) that attended rapid VI responding. On the other hand, when poor (a circumstance arranged by restricting the number of responses per session) a monkey could optimize its daily food stock by responding more slowly, because the lower the VI response rate, the fewer the responses per reinforcer. Although slow responding increases response efficiency by reducing the number of responses per reinforcer, it is not efficiency per se that

would be presumed to cause reduced rates. Rather, it is the fact that when the number of responses per session is restricted, reduced response rates lead to more food per session than higher response rates do.

METHOD

Subjects

Two adult male rhesus monkeys (Monkeys 5 and 6) began the experiment at their free-feeding weights. In general, they received no food other than that scheduled by the experimental procedures. There were two exceptions: (a) a multivitamin tablet was given three times per week, and (b) during any day that a session was not conducted, 150 food pellets were given. Therefore, their body weights were allowed to vary with experimental conditions. Both monkeys had long histories of responding on single and concurrent schedules of reinforcement.

Apparatus

The monkeys were housed individually in cages measuring 70 cm wide, 70 cm deep, and 80 cm high. A panel accessible between the cage bars was attached to one side of the cage. It contained six keys (4.5 cm diameter) arranged in two horizontal rows of three (6.5 cm apart horizontally and 11 cm apart vertically, center to center). A circular opening (5.5 cm diameter), located to the left of the keys, permitted access to a hopper where food pellets were delivered.

Procedure

Each daily session began with the illumination, with white light, of the center key in the top row of keys. Responses on this key were reinforced with the delivery of four 750-mg banana-flavored Noyes monkey pellets according to a VI 3-min schedule, the intervals of which were determined according to Flesher and Hoffman's (1962) progression. During the reinforcement cycle, pellets dropped into the hopper at the rate of one per second. The VI schedule resumed operation immediately after the last pellet was released by the food dispenser.

In Phase 1, sessions ended after 3 hr. In Phase 2, sessions ended after subjects emitted 20% of the average number of responses per session during the last five sessions of Phase

Table 1

Experiment 1. Summary of experimental conditions, response rates, and number of responses per reinforcer. Results are means of the last six sessions, with standard deviations in parentheses. Response rates are based on the first 60 min of each session.

Mon-key	Phase	Responses per session	Number of sessions	Responses per minute	Responses per reinforcer
5	1	unlimited	33	36	104 (23)
	2	1,202	15	29	116 (49)
	3	3,000	14	77	173 (31)
	4	unlimited	12	48	128 (11)
6	1	unlimited	34	20	43 (10)
	2	395	10	28	129 (123)
	3	1,500	15	57	177 (72)
	4	unlimited	12	19	64 (22)

1. Because this restriction reduced food intake to unacceptable levels, this fixed number of responses was increased in Phase 3 from 1,202 to 3,000 for Monkey 5 (50% of Phase 1 response levels), and from 395 to 1,500 for Monkey 6 (76% of Phase 1 levels). In Phase 4, sessions once again lasted for 3 hr, with no restriction on the number of responses.

Conditions ended when the number of responses per reinforcer (Figure 1) was judged to show no substantial changes in trend or pattern over six sessions. These stability assessments began after 10 sessions were completed. Table 1 presents, for each subject in each phase, the number of responses permitted per session and the number of sessions defining a phase.

RESULTS

In the present experiment, developing a meaningful summary measure of response rate was problematic because in Phases 1 and 4, the 3-hr session duration substantially exceeded the period of time during which subjects responded. In consequence, the overall rate could be low not because monkeys responded slowly, but rather because the period from the cessation of responding to session's end inflated the denominator used in determining response rate. To minimize this factor's contribution to the measure of response rate, all Phase 1 and Phase 4 response rates in Table 1 are based on responding during the first hour of the session—a period during which both monkeys were responding consistently. In Phases 2 and




Fig. 1. Responses per reinforcer as a function of days for Monkeys 5 and 6 in Experiment 1. Spaces without data points refer to days when the experiment was not conducted or when there was an equipment or data-collection failure.

3, in which sessions were much shorter, response rates were calculated in the usual way, by dividing total number of responses by the session time. Except for Monkey 5, Table 1

shows that response rates were higher when responses per session were restricted than when they were not.

Figure 2 presents for both subjects the num-

ber of responses (top panels), reinforcers (middle panels), and responses per reinforcer (bottom panels) during low- and high-wealth conditions. These data mirror those of Table 1: Except for Phase 2 for Monkey 5, the number of responses per reinforcer was higher when responding was restricted (wealth was low) than when it was not restricted (wealth was high).

Figure 1 presents the number of responses per reinforcer for each subject throughout the experiment. (One-day lapses in data presentation typically indicate an equipment failure. The several-day break starting with Day 11 was for tuberculosis testing.) These data demonstrate that when responding was restricted (Phases 2 and 3), not only did the number of responses per reinforcer tend to increase on average from the unrestricted phases (Phases 1 and 4), but so too did within-phase variability. This variability effect is represented by the mean and standard deviation of the number of responses per reinforcer during the last six sessions of each phase of the experiment (see Table 1).

DISCUSSION

On VI schedules, the probability of reinforcement increases with increasing interresponse times (IRTs). If the number of responses permitted in a session is reduced, an organism can defend its within-session food intake by increasing the average IRT—in other words, by responding more slowly. Such a pattern increases the likelihood that a response will result in reinforcement, and, unless the response restriction is severe, can compensate for the decline in total reinforcement that would accompany an unchanged response rate.

There was little evidence to support the idea that monkeys would optimize the stock of food by responding more slowly (Figure 2). In fact, responding generally became *less* efficient as the number of responses per session was restricted. These results can be explained in terms of a behavior-analytic notion of response strength advocated by Vaughan and Miller (1984). According to this account, the present study's response-restriction conditions increased the reinforcing value of food by reducing food intake and lowering body weight. Because a strength account posits that behavior output is correlated with reinforcer value, the account is consistent with the present study's




Fig. 2. Mean daily responses, reinforcers, and responses per reinforcer during the last six sessions of high- and low-wealth conditions for Monkeys 5 and 6 in Experiment 1.

major finding—that the number of responses per reinforcer increased when the number of responses per session was restricted.

A second explanation attributes response-rate increases to consumption restrictions in

Phases 2 and 3. According to this view, response rates increased because the monkeys were deprived of food to a greater degree, and the deprivation energized action. This account is simpler than the response-strength explanation. In that account, deprivation did not directly change behavioral output. Rather, it changed reinforcer value, and then the changed value of the reinforcer changed the rate of responding.

The present report cannot answer whether the drop in food consumption increased response rates indirectly by increasing the reinforcing strength of food, or increased response rates directly. However, neither possibility compromises the claim that this study fails to show stock optimizing.

This failure to optimize stocks may reflect monkeys' insensitivity to the method used to restrict responding rather than their insensitivity to restriction *per se*. The wealth-restricting contingency—reducing the number of responses that could be emitted in a session—could be experienced only once per day, and offered no obvious means for response number to function as a cue. The restriction of responding may have been too difficult for monkeys to discriminate, enabling nonoptimizing processes to govern behavior.

This interpretation can be tested experimentally if the number of available responses is made more explicit. One way to do this is to replace responses with tokens—that is, define the operant as the deposit of a token in a slot. By controlling the number of tokens, a monkey can be made either rich or poor, and it can visually monitor the unspent tokens. When wealthy, it can afford to deposit tokens rapidly and to maximize reinforcer immediacy (Silberberg *et al.*, 1988). However, when poor, it should husband its limited cash by depositing tokens more slowly. The longer IRTs produced by a lowered rate would raise the probability that a token deposit would produce VI reinforcement and would compensate for the reduced number of tokens by functionally lowering the price of the reinforcer (deposits per pellet).

EXPERIMENT 2

The second experiment was similar to the first, except that wealth levels were manipulated by adjusting the number of token deposits

per session rather than the number of key presses. In the high-wealth conditions, tokens deposited were immediately returned to the monkey, so that wealth was unlimited. In the low-wealth conditions, deposited tokens were not returned, and the allotted number of tokens, if deposited at the same rate as in the high-wealth conditions, would result in reductions in food intake and body weight.

METHOD

Subjects

Two adult male rhesus monkeys (Monkeys 1 and 2), both experimentally naive, served as subjects. They received no food other than that scheduled by the experimental procedures, except a multivitamin tablet three times per week. They began the experiment at their free-feeding weights of 10 kg each.

Apparatus

Each subject was housed in a cage measuring 70 cm wide, 70 cm deep, and 80 cm high. Two aluminum panels (61 cm wide, 43 cm high), one containing a reservoir of tokens and a slot through which they could be deposited, and another containing a cup into which food pellets could be delivered, were attached to opposite side walls of each subject's cage. A Plexiglas vertical deposit slot (1 cm by 4.4 cm) was located 16.5 cm from the top of the token panel and 12 cm from its right edge. A hole (7 cm diameter), located 17 cm below the lower edge of the token slot, provided access to a supply of black plastic tokens (3.8 cm diameter, 0.16 cm thick). A hole (5 cm long, 3 cm high) in the reinforcer panel permitted access to a recessed food cup.

Procedure

Both subjects, deprived of food for approximately 22 hr, were trained by shaping to remove a token from the supply and deposit it through the slot in the token panel. Each token deposit caused two food pellets to be delivered into the cup in the reinforcer panel. Once responding seemed reliable, this continuous reinforcement schedule was replaced by a VI 3-min schedule that delivered four food pellets for the first token deposit following the end of an interval. The arrangements for scheduling and delivering the pellets were the same as in Experiment 1.

Table 2

Experiment 2. Summary of experimental conditions, and number of responses per reinforcer. Results are means of the last six sessions, with standard deviations in parentheses.

Monkey	Phase	Wealth	Number of sessions	Responses per reinforcer
1	1	high	26	14 (0.9)
	2	low	18	5 (0.8)
	3	high	15	7.5 (0.5)
	4	low	24	5.9 (1.8)
	5	high	24	10.8 (0.6)
2	1	high	26	6.4 (0.4)
	2	low	24	7.3 (3.7)
	3	high	15	10.2 (0.7)
	4	low	26	5.6 (1.1)
	5	high	24	17.5 (1.0)

In all phases of the study, 150 tokens were placed in an open reservoir in the response panel. When wealth was high, any tokens deposited through the token slot were returned to the reservoir, ensuring that subjects had essentially an infinite supply of tokens. When wealth was low, deposited tokens were not returned, limiting monkeys to a maximum of 150 token deposits per session.

A white light located behind the Plexiglas token slot was illuminated when the session began and extinguished at session's end. All sessions lasted 3 hr or until no tokens were left, whichever occurred first.

For both subjects, high- and low-wealth conditions alternated through the five phases of this experiment. Each phase ended when the number of responses per reinforcer was judged stable over the most recent six sessions (except for Phase 2 for Monkey 2). The order of conditions and the number of sessions per condition are presented in Table 2.

RESULTS

Figure 3 presents for both subjects the mean daily number of token deposits (top panels), reinforcers (middle panels), and token deposits per reinforcer (bottom panels) during the session at the end of high-wealth and low-wealth conditions. The numbers atop bars in the top panel identify monkeys' body weights (in kilograms) before one of the last six sessions of a particular condition. The session selected for weighing was at the experimenter's convenience and varied from condition to condition.




Fig. 3. Mean daily responses, reinforcers, and responses per reinforcer during the last six sessions of high- and low-wealth conditions for Monkeys 1 and 2 in Experiment 2.




Fig. 4. Responses per reinforcer as a function of days for Monkeys 1 and 2 in Experiment 2.

Both monkeys responded more when wealthy than when poor; however, they responded more efficiently when poor, a fact represented by the finding that the responses-per-reinforcer ratio was consistently lower during

low-wealth conditions than during high-wealth conditions.

Figure 4 shows how the number of responses per reinforcer changed throughout Experiment 2. Generally speaking, the levels




Fig. 5. Mean of the median daily response rates for Monkeys 1 and 2 during the last six sessions of each phase of Experiment 2. The numbers atop each bar are the standard deviations. See text for details.

of within-phase variability in responses per reinforcer were similar across the low-wealth and high-wealth conditions.

Figure 5 presents the mean of each monkey's daily median response rate during the last six sessions of each condition. Each median was calculated by recording IRTs in 1-s classes; the reciprocal of the median IRT defined the median rate for the session. Median rates (rather than mean rates) were used because occasionally monkeys quit responding before session's end, especially during high-wealth conditions. Although median rates are sensitive to response rates while the monkey is responding, they are insensitive to when the subject quits responding in a session. Thus, in this context, median rates are a superior measure of responding than mean rates. The value in Figure 5 is the mean of these median rates, and the value atop each bar is the standard deviation. The monkeys responded more slowly when poor than when rich. Finally, whether rich or poor, these token deposit rates were

much lower than the rates in Experiment 1, in which a key press served as the operant.

DISCUSSION

In Experiment 1, optimizing food intake required that VI response rates decrease when the number of responses permitted in a session was reduced. Instead, monkeys responded more rapidly, accentuating the size of their food loss from the levels that would have been obtained had rates remained unchanged. This result violates the predictions of a stock-optimizing account of VI performance.

In Experiment 2, we entertained an explanation of this finding that preserved a role for stock optimizing in VI performance. According to this explanation, subjects in Experiment 1 would have optimized, lowering response rates when the number of allotted responses per session was reduced, if the relation between response output and session's end had been more discriminable. To test this thesis, we gave monkeys a supply of tokens that either

diminished with deposits or did not. We assumed that both the size of the token supply and the effect of token deposits on that size could be readily discriminated. When we repeated Experiment 1's design with tokens instead of key presses, monkeys responded more slowly when the number of token deposits per session was restricted. This outcome is consistent with our discrimination hypothesis and the operation of stock optimizing in VI responding.

Despite conformity between our data and the predictions of stock optimizing, the present results do not establish that stock optimizing governs behavior. Other explanations may be viable. For example, these results can possibly be explained in terms of conditioned reinforcement principles. According to such an account, the association of food reinforcement with tokens may have made the tokens conditioned reinforcers. If the tokens were reinforcing, their removal in the restricted-token conditions may have been aversive, leading to reduction of the rate of depositing the tokens. The credibility of such a conditioned reinforcement interpretation awaits further work.

GENERAL DISCUSSION

Humans make economic decisions based on both their rate of pay and the total pay received. To illustrate, imagine that several graduate students were offered two alternatives for summer employment. Job 1 pays \$2,000 for 1 month's work, whereas Job 2 pays \$3,000 for 3 months' work. We would not expect unanimity of opinion in job selection. Those optimizing rate (flow) of pay would opt for Job 1, whereas those more interested in optimizing aggregate pay (stock) would choose Job 2.

Despite the operation of both stock and flow variables in their own economic decision making, behavior analysts typically attribute behavioral control to reinforcer flows rather than to reinforcer stocks. One can defend this emphasis by noting that there is little evidence to show that reinforcer stocks play an important role in the control of behavior. On the other hand, there are also few experimental tests of what that role might be (for some examples, see Collier *et al.*, 1977; Hursh, 1980; Jensen & Rey, 1968).

The paucity of data on control by reinforcer

stocks may be due to the longstanding practice of maintaining constant body weights in subjects whose behavior is reinforced by presentation of food. As long as body weight remains fixed, it is not possible to define a functional relation between food-stock size and behavior. Hence, it may be a mistake to interpret the absence of many demonstrations of the role of reinforcer stocks on behavior as reflecting the fact that this variable is of little consequence.

The predisposition to evaluate reinforcers in terms of their flow rather than their stock is reflected in how behavior analysts and optimal foraging theorists model behavior (e.g., Caraco *et al.*, 1980; Herrnstein, 1961; Herrnstein & Heyman, 1979). For example, Herrnstein's (1961) matching model has parameters to accommodate changes in the rate of reinforcement, but none to accommodate changes in aggregate reinforcement. One exception is Hursh's (1980) position that ensuring a dependence between behavioral output and aggregate reinforcement—a circumstance usually arranged by requiring subjects to earn their entire daily food ratio within the experimental session—results in behavior that defends food intake (see also Collier, Hirsch, & Hamlin, 1972; Hirsch & Collier, 1974). This behavior in a "closed" economy is, of course, a form of stock optimization.

Time is a critical feature in virtually all behavior-analytic models. As any student of schedule effects recognizes, this criticality is well placed (see Ferster & Skinner, 1957). Nevertheless, time has a prominence in behavior analysis that is unmatched in microeconomic models of behavior. One reason for this difference is that economists have found that their models often sacrifice little in accuracy by ignoring the effects of time; when time is considered, such as in calculating the present value of alternative outcomes, it must be represented in years to influence a model's predictions to a significant degree.

Certainly the situation differs in research with nonhumans, where we often find a few seconds' delay in one reinforcer versus another can powerfully influence action (Rachlin & Green, 1972). Although this apparent phyletic difference may suggest why flow variables are of greater importance to the behavior analyst than to the economist, it does not endorse the behavior-analytic tradition of interpreting reinforcement effects solely in terms of flow vari-

ables, such as rate of reinforcement. Indeed, as the present study makes clear, nonhumans are sensitive not only to the temporal characteristics of reinforcement but also to the aggregate amount of reinforcement produced.

REFERENCES

- Caraco, T., Martindale, S., & Whittam, T. S. (1980). An empirical demonstration of risk-sensitive foraging preferences. *Animal Behaviour*, **28**, 820-830.
- Collier, G., Hirsch, E., & Hamlin, P. (1972). The ecological determinants of reinforcement in the rat. *Physiology and Behavior*, **9**, 705-716.
- Collier, G., Hirsch, E., & Kanarek, R. (1977). The operant revisited. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 28-52). Englewood Cliffs, NJ: Prentice-Hall.
- Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.
- Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, **5**, 529-530.
- Hamm, S. L., & Shettleworth, S. J. (1987). Risk aversion in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **13**, 376-383.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, **4**, 267-272.
- Herrnstein, R. J. (1990). Rational choice theory: Necessary but not sufficient. *American Psychologist*, **45**, 356-367.
- Herrnstein, R. J., & Heyman, G. M. (1979). Is matching compatible with reinforcement maximization on concurrent variable interval, variable ratio? *Journal of the Experimental Analysis of Behavior*, **31**, 209-223.
- Hirsch, E., & Collier, G. (1974). Effort as determinant of intake and patterns of drinking in the guinea pig. *Physiology and Behavior*, **12**, 647-655.
- Hursh, S. R. (1980). Economic concepts for the analysis of behavior. *Journal of the Experimental Analysis of Behavior*, **34**, 219-238.
- Jensen, G. D., & Rey, R. P. (1968). Runway performance under "horn of plenty" conditions versus gradual diminution of reward supply. *Journal of Experimental Psychology*, **76**, 7-11.
- Lea, S. E. G., Tarpy, R. M., & Webley, P. (1987). *The individual in the economy*. Cambridge, UK: Cambridge University Press.
- Logue, A. W., Chavarro, A., Rachlin, H., & Reeder, R. W. (1988). Impulsiveness in pigeons living in the experimental chamber. *Animal Learning & Behavior*, **16**, 31-39.
- Mazur, J. E. (1988). Choice between small certain and large uncertain reinforcers. *Animal Learning & Behavior*, **16**, 199-205.
- Rachlin, H., Battalio, R., Kagel, J., & Green, L. (1981). Maximization theory in behavioral psychology. *The Behavioral and Brain Sciences*, **4**, 371-417.
- Rachlin, H., & Green, L. (1972). Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior*, **17**, 15-22.
- Sakagami, T., Hursh, S. R., Christensen, J., & Silberberg, A. (1989). Income maximizing in concurrent interval-ratio schedules. *Journal of the Experimental Analysis of Behavior*, **52**, 41-46.
- Samuelson, P. A., & Nordhaus, W. D. (1985). *Economics* (12th ed.). New York: McGraw-Hill.
- Shurtleff, D., & Silberberg, A. (1990). Income maximizing on concurrent ratio-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, **53**, 273-284.
- Silberberg, A., Warren-Boulton, F. R., & Asano, T. (1988). Maximizing present value: A model to explain why moderate response rates obtain on variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, **49**, 331-338.
- Timberlake, W. (1984). A temporal limit on the effect of future food on current performance in an analogue of foraging and welfare. *Journal of the Experimental Analysis of Behavior*, **41**, 117-124.
- Timberlake, W., Gawley, D., & Lucas, G. (1987). Time horizons in rats foraging for food in temporally separated patches. *Journal of Experimental Psychology: Animal Behavior Processes*, **13**, 302-309.
- Vaughan, W., Jr., & Miller, H. L., Jr. (1984). Optimization versus response strength accounts of behavior. *Journal of the Experimental Analysis of Behavior*, **42**, 337-348.

Received August 3, 1991

Final acceptance November 17, 1992